Mon. Not. R. Astron. Soc. **000**, 1–10 (2015) Printed 12 April 2016 (MN LATEX style file v2.2)

Characterising the rotational irregularities of the Vela pulsar from 21 yr of phase-coherent timing

R. M. Shannon⋆[1,2], L. T. Lentati[3], M. Kerr[1], S. Johnston[1], G. Hobbs[1], & R. N. Manchester[1]

[1] *CSIRO Astronomy and Space Science, Australia Telescope National Facility, Box 76 Epping, NSW, 1710, Australia.*
[2] *International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia.*
[3] *Astrophysics Group, Cavendish Laboratory, JJ Thomson Avenue, Cambridge CB3 0HE, UK.*

arXiv:1604.02661v1 [astro-ph.HE] 10 Apr 2016

Accepted XX ; in original form 12 April 2016

ABSTRACT

Pulsars show two classes of rotational irregularities that can be used to understand neutron-star interiors and magnetospheres: glitches and timing noise. Here we present an analysis of the Vela pulsar spanning nearly 21 yr of observation and including 8 glitches. We identify the relative pulse number of all of the observations between glitches, with the only pulse-number ambiguities existing over glitch events. We use the phase coherence of the timing solution to simultaneously model the timing noise and glitches in a Bayesian framework, allowing us to select preferred models for both. We find the glitches can be described using only permanent and transient changes in spin frequency, i.e., no step changes in frequency derivative. For all of the glitches, we only need two exponentially decaying changes in spin frequency to model the transient components. In contrast to previous studies, we find that the dominant transient components decay on a common $\approx$ 1300 d time scale, and that a larger fraction ($\gtrsim$ 25%) of glitch amplitudes are associated with these transient components. We also detect shorter-duration transient components of $\approx$ 25 d, as previously observed, but are limited in sensitivity to events with shorter durations by the cadence of our observations. The timing noise is well described by a steep power-law process that is independent of the glitches and subdominant to the glitch recovery. The braking index is constrained to be < 8 with 95% confidence. This methodology can be used to robustly measure the properties of glitches and timing noise in other pulsars.

Key words: pulsars: general – pulsars: specific (PSR B0833−45) – stars: neutron

1 INTRODUCTION

Pulsars are celebrated for the predictability of the arrival times of their pulses. The power of the pulsar-timing method is realised when a phase-coherent model of the pulse times of arrival (TOAs) is achieved; that is, when there is a solution that unambiguously accounts for every rotation of the pulsar. For example, in an analysis of 25 yr of observations of the first millisecond pulsar, PSR B1937+21, the last observed TOA in 2010 is 518,595,058,289 periods after first TOA in 1986 (Shannon et al. 2013). Phase coherence enables the detection of subtle effects that only slightly modify the pulse frequency to accumulate in the arriving phase of pulses. These effects include variations in the orbit of binary pulsars associated with general-relativistic effects and potentially the passage of gravitational waves with frequencies in the nanohertz range. However, nearly every pulsar shows evidence for intrinsic spin irregularities which also alter the TOAs. Irregularities are phenomenologically bifurcated into two forms, timing noise and glitches, with glitches more common in younger pulsars.

Timing noise manifests as a red-noise (time-correlated) process in the TOAs, and is typically described by a wide-sense stationary stochastic process (Groth 1975), modelled as either random walks in the pulsar spin parameters, or with a more general power spectrum. The origin of this noise is unclear. While some of it may be caused by variable torques associated with changes in the pulsar magnetosphere state (Lyne et al. 2010), a significant fraction is likely associated with rotational irregularities interior to the star (Melatos & Link 2014). The properties of timing noise vary markedly across the pulsar population, with its

⋆ E-mail: ryan.shannon@csiro.au

© 2015 RAS

strength depending on the pulsar spin frequency ν and frequency derivative $\dot{\nu}$ (Shannon & Cordes 2010), and likely other factors because of the large dispersion levels of timing noise between pulsars. For example, the most stable millisecond pulsar (MSP), PSR J1909−3744 shows no evidence for instabilities with phase variations limited $\lesssim 100$ ns ($\approx 0.3\%$ of pulse phase) over 11 yr (Shannon et al. 2015). In contrast, the timing noise in young pulsars and magnetars can contribute many cycles of pulse phase on week to month time scales (Livingstone et al. 2011), both making it difficult to find a phase coherent solutions in poorly sampled data and presenting challenges to TOA-modelling algorithms. While finding the origin of timing noise is important for understanding neutron stars, it is also necessary to account for timing noise as part of a general timing model for the pulsar. This is necessary to eliminate, or at least mitigate, bias in the estimation of other parameters in the model (Coles et al. 2011; van Haasteren & Levin 2013; Lentati et al. 2014; Kerr et al. 2015).

In addition to exhibiting timing noise, pulsars can experience glitch events (e.g., Espinoza et al. 2011; Yu et al. 2013), in which they are observed to suddenly change spin state, with the most significant component being increases in spin frequency that can exceed $1:10^5$ (Manchester & Hobbs 2011). Glitches have been modelled with permanent changes in ν and $\dot{\nu}$, as well as transient components, in which changes in ν are modelled to decay (typically exponentially) on time scales τ. Occasionally multiple glitch-decay components (with different time scales) are invoked, particularly when high cadence (daily or higher) observations (e.g., Dodson, McCulloch & Lewis 2002). Glitch components are often identified and characterised by searching for variation of spin frequency and frequency derivative in subsets of the data (Lyne et al. 2015). This method is suboptimal because it does not utilise the phase coherence of the pulse arrival times. The study of glitches is further complicated by the presence of timing noise.

There are two prevailing theoretical models for glitches. In the first, glitches are associated with the transfer of angular momentum between the superfluid interior and solid crust of the neutron star (Anderson & Itoh 1975; Alpar et al. 1984). A superfluid component is present in the core of the neutron star and a portion pervades the inner crust (Baym, Pethick & Pines 1969). The quantised angular velocity vortices can pin on the nuclear lattice of the crust. As the crust and the normal fluid component of the NS spin down because of electromagnetic braking, differential angular momentum is built up between the pinned vortices and the other components. Eventually an external trigger or the Magnus force (e.g., Melatos & Warszawski 2009) causes the vortices to unpin and transfer angular momentum to the crust, spinning the star up. In the second model, glitches are associated with star-quakes in the solid, crystalline crust of the neutron star. The quakes are the result of changes in the equilibrium configuration as the oblate star relaxes toward a spherical state as it slows down and cools (Ruderman 1969). The latter model is presently disfavoured for most pulsars because it cannot account for the amplitude distribution and event rate observed in glitches (Haskell & Melatos 2015), though it might be suitable for a few pulsars such as the Crab pulsar. Empirical evidence suggests that glitches can be modelled using a deterministic signal in the pulsar timing model (Edwards, Hobbs & Manchester 2006) with only a few parameters.

Theoretical models predict that glitch decay time scales should be constant for individual pulsars if the underlying physics driving glitch recovery is a linear process and that the transient components should have comparable magnitudes to the permanent components (Haskell & Melatos 2015). However, the transient components hitherto measured have a small contribution relative to the permanent component and have variable decay time scales (Espinoza et al. 2011).

Many young pulsars also show evidence for measurable braking that is attributed to pulsar spin down. This is a deterministic process that primarily manifests as a second derivative of spin frequency $\ddot{\nu}$, and is parametrized by a braking index n (where $\dot{\nu} \propto \nu^n$) . For electromagnetic braking associated with a dipolar magnetic field $n = 3$; however measured braking indices often depart from this markedly, with contributions likely arising from angular momentum loss from particle winds and free precession.

Here we study the timing properties of the Vela pulsar (Large, Vaughan & Mills 1968), a relatively young (characteristic age of $\tau_c \approx 11$ kyr) pulsar that shows both large levels of timing noise and glitches at quasi-regular intervals. The Vela pulsar is the brightest known pulsar at decimetre wavelengths, with period-averaged flux density of $\sim$ 1 Jy. The first glitch in any pulsar was identified in the Vela pulsar (Radhakrishnan & Manchester 1969; Reichley & Downs 1969) and 16 subsequent glitches have been detected (Yu et al. 2013). Timing noise analysis of the pulsar (Cordes & Helfand 1980; Cordes, Downs & Krause-Polstorff 1988) has typically been constrained to intervals between glitches. Glitch analyses similarly have been rarely conducted including the effects of timing noise. Similarly, the observation of a very low braking index of $n = 1.4 \pm 0.2$ reported for the pulsar (Lyne et al. 1996) did not account for timing noise.

To analyse both the glitches and the timing noise simultaneously, we use a timing solution for the Vela pulsar that spans 21 yr of observation and 8 glitch events, enabling us to examine the stationarity of the timing noise and robustly estimate glitch parameters. In section 2, we present the observational data used. In section 3, we discuss the timing analysis and Bayesian methodology employed. In section 4, we compare phenomenological models for the timing noise and the glitches and select a preferred model. In section 5, we discuss the implications of this model. In section 6, we give our conclusions.

2 OBSERVATIONS

Our highly heterogeneous data set comprises 1231 TOAs obtained with the 64-metre Parkes radio telescope between 1992 December 12 and 2014 January 14. While observations were conducted at frequencies between 0.4 and 23 GHz, most were made at a central frequency of ≈ 1.4 GHz. Prior to 2003 observations were made with a series of analogue-filterbank and digital-autocorrelation spectrometers; these observations are described in detail in Wang et al. (2000) and Yu et al. (2013). Most recently, the pulsar has been observed with digital polyphase filterbank spectrometers

as part of a programme to monitor pulsars of interest to the *Fermi* gamma-ray observatory (Weltevrede et al. 2010). These observations have monthly cadence with a central observing frequency close to 1.4 GHz and semi-annual cadence with a dual-band system capable of observing simultaneously at central frequencies of 0.73 and 3.0 GHz.

The primary data in this analysis are TOAs, formed by correlating observations that have been averaged in frequency, time and (where recorded) polarisation with a template, using the commonly applied Fourier phase gradient method, described in Taylor (1992), and implemented in the pulsar analysis code PSRCHIVE (Hotan, van Straten & Manchester 2004). Templates were produced individually for each backend/observing-band combination using an analytic model fitted to the average profile from that combination. Offsets between the backends were included in the timing model, as discussed in Section 3. The cross-correlation method assumes that the data can be described by the template and additive white noise. For our observations, this is not the case. Distortions of the pulse profile, especially prevalent in older observations, are introduced both by the high flux density (in excess of system equivalent flux density) of the pulsar, and the large dispersion sweep of the pulsar relative to the pulse phase and frequency resolution of the observations.

Saturation of the amplifiers, other non-linear effects in the receiver and downconversion chain, and low-bit digitisation can lead to artefacts in the pulse profile, such as apparent negative flux density on the leading and trailing edges of the pulse (Jenet et al. 1998). Older observations were recorded with analogue-filterbank spectrometers with single-bit digitisers and were especially susceptible to these artefacts. Additionally, the pulse profile can be artificially broadened if the dispersive delay across an individual channel bandwidth is larger than the pulse-phase resolution of the observation. Given the relatively narrow pulse (2.1 ms), and relatively high dispersion measure ($68\,\mathrm{pc\,cm^{-3}}$), older observations conducted with wide channels at low frequency show this type of broadening. Even in more recent observations where instrumental effects are minimised, stochasticity in the pulse shape introduces additional timing error (referred to as pulse jitter, Cordes & Downs 1985) that limits the timing precision of the observations. The effects of all of these distortions are secondary to TOA variations induced by timing noise and glitch events. It is however necessary to account for these effects in the analysis, in particular when modelling transient glitch components in our sparsely sampled data set. While we do not account for them while measuring TOAs (Lentati & Shannon 2015), we account for their effect in the pulsar timing model, as discussed in the next section.

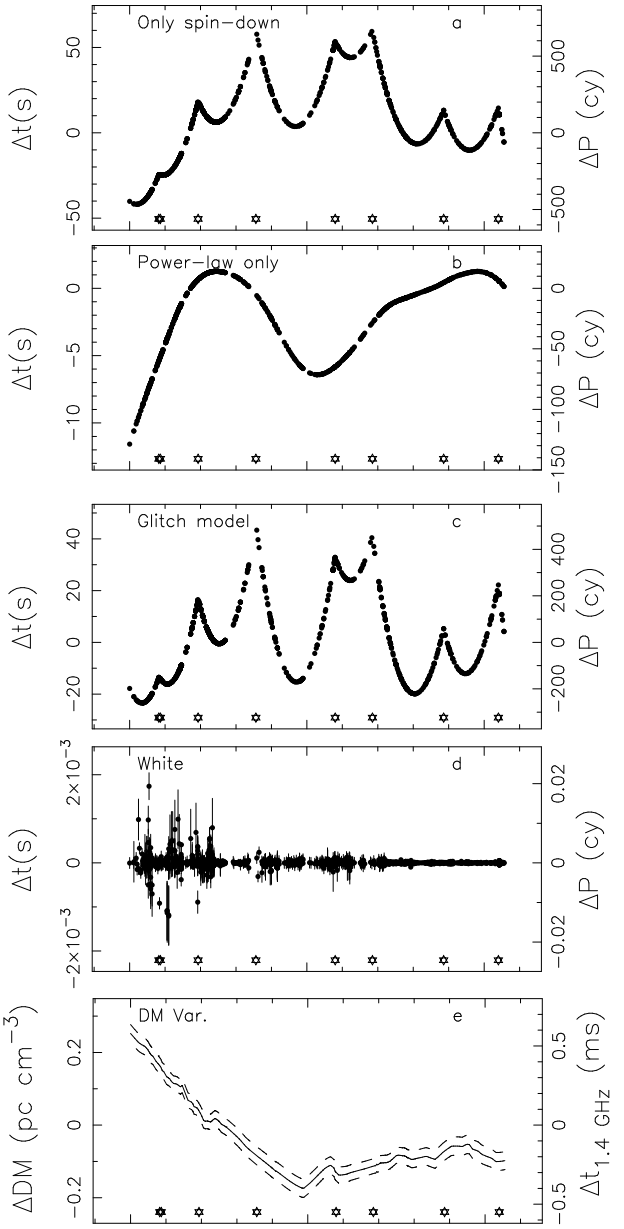

Figure 1. Residual arrival times for maximum-likelihood models of the Vela pulsar, measured in units of time Δt and cycles of phase ΔP. The stars show the epochs of the glitches. *a*: Only fitting for the spin frequency and frequency derivative. *b*. Fitting for the glitches but assuming power-law noise. *c*: Modelled glitch signal from solution presented in panel b. *d*: Whitened residuals for maximum-likelihood solution. *e*: Dispersion measure variations for the maximum-likelihood solution (solid line) . The dashed line shows the 1σ uncertainties on the realisation.

3 TIMING ANALYSIS

The presence of strong timing noise and glitch events make it difficult to produce phase-connected solutions over long data spans for young, energetic pulsars like the Vela pulsar. As the data spans increase, the amplitude of the timing-noise signal increases rapidly (with the timing noise having a power spectral density $P_r(f) \propto f^{-5\pm2}$ across the population, Shannon & Cordes 2010), and relative to a spin period and period derivative at some fiducial epoch, the arrival times diverge. Even if an initial phase-connected solution exists, it typically has hitherto been difficult to fit the solution because small changes in model parameters can change the residual arrival times by more than one cycle of pulse phase.

Within the pulsar-timing code TEMPO2 (Hobbs, Edwards & Manchester 2006), it is now possible to use the relative pulse numbers of the TOAs as a reference when measuring the goodness of fit or the

likelihood of a model. If the level of timing noise is large so that phase connecting the TOAs over the entire data set is difficult or impossible, the relative pulse numbers for the entire data set can be determined after producing phase-connected time solutions for subsets of observations. The overlapping intervals are chosen to be sufficiently long that the relative pulse numbering can be checked for agreement These solutions do not need to be physical, so timing noise can be whitened using sinusoids (Hobbs et al. 2005) or derivatives of the pulse spin frequency. The consistency of the solutions can be checked using TOAs common to different intervals, i.e., relative pulse numbers should be the same for overlapping observations.

Ambiguities in the pulse numbering potentially exist immediately following glitch events. If the epoch of the glitch is poorly constrained, the glitch amplitude is large and the observation cadence is poor around the glitch epoch, there may be uncertainty in the number of rotations of the pulsar between the glitch and the first post-glitch observation. Possible glitch epochs, corresponding to changes in the pulse numbering by one unit, are separated by $\Delta\nu_g^{-1}$, where $\Delta\nu_g$ is the change in spin frequency associated with the glitch. For most of the glitches in our data set, the glitch epochs have been previously reported to sufficient accuracy such that we can unambiguously identify the pulse rotations through the glitch event. For these glitches, the maximum-likelihood glitch phase parameter is found to be within < 0.1 cycles of zero offset. If the glitch epoch were perfectly determined, the glitch phase parameter would be zero. However, a few of the glitches epochs are poorly constrained. In this case we are unsure of the relative pulse numbering through the glitches. We tested the effects of having incorrect pulse numbers through the glitches, and found that no parameters in our analysis (beyond the glitch phase) were significantly affected, because the glitch epochs are constrained to a time much shorter than the fastest varying process in our model. This is to be expected, because we have included the glitch phase parameter in our timing analysis. Because we analytically marginalise over the glitch phase parameter, we are marginalising over the pulse-number uncertainty that exists at glitch epochs.

Between glitches, we confirmed our pulse numbering is correct by whitening the entire dataset with a series of sinusoids (Hobbs et al. 2005). Incorrect pulse numbering would result in TOAs that have a random phase offset from the other arrival times, and residuals distributed across a full cycle of pulse phase in the whitened solution. We used a series of 40 sinusoids with periods ranging from 23.2 to 0.58 yr. This was sufficient to model most of the red timing noise with the residual TOAs constrained to ≈ 0.05 of pulse phase. We found no outlying points. In contrast, when we purposely introduced pulse numbering errors, outlying TOAs were easily identified.

Using this timing solution, we can directly calculate the absolute residual pulse phase of the TOAs to the model, even if the difference between the model and data are $\gg 1$ cycle of phase. We used the code TEMPONEST (Lentati et al. 2014) to construct a complete timing model from the phase-connected solution. The Bayesian framework implemented in TEMPONEST enables us to simultaneously model stochastic parameters (e.g., timing noise) and deterministic parameters (e.g., glitch parameters) of interest, and marginalise over nuisance parameters of no interest to the analysis. For example we were able to search over a wide range of glitch decay times and glitch epochs, which is not possible with the fitting algorithm internal to TEMPO2.

This approach also enabled us to select favoured models through the use of Bayesian evidence (the integral of the likelihood over the parameter space weighted by the prior), which can be used to calculate the relative probabilities of different models. TEMPONEST was recently enhanced to use the PolyChord algorithm (Handley, Hobson & Lasenby 2015) to sample posterior distributions and calculate the evidence, enabling more efficient searches and more robust calculation of evidences in high-dimensional parameter spaces than the alternative MULTINEST algorithm (Feroz, Hobson & Bridges 2009).

We also had to modify the code to incorporate higher precision (128-bit and 256-bit) floating-point arithmetic[1], because of the dynamic range required to model the arrival times. The highest precision measurements have individual TOA errors (after accounting for systematics and pulse shape variations) of a few μs, while the plausible red-noise variations over the data span exceed 1000 s. This corresponds to a dynamic range of 10^{18} for the noise covariance. Matrices with this dynamic range need to be inverted as part of our analysis. For our analysis, we found that 128-bit precision was sufficient, and the 256-bit precision was notably slower.

Our model for the TOAs includes deterministic variations (Edwards, Hobbs & Manchester 2006) to account for the pulsar spin-down and astrometric terms. There are position, proper motion, and parallax measurements from long-baseline radio interferometry that exceed the precision we can obtain through pulsar timing by factors of $\gg 10$ (Dodson et al. 2003) and are consistent with our measurements. We therefore held the value fixed at the interferometrically determined position in our analysis[2]. We analytically marginalised over deterministic parameters that are linear (or linearisable) in the timing model (van Haasteren et al. 2009), greatly reducing the time required to estimate non-linear and stochastic components of the timing model.

We also included terms to account for stochastic time-independent (white-noise) and time-correlated (red-noise) contributions to TOAs. In addition to the white noise associated with the formal TOA uncertainty, terms are included to account for intrinsic shape variations and instrumental distortions, and other instrumental and astrophysical effects that are temporally uncorrelated between observations. We model the white noise by adjusting the uncertainty on individual TOAs to be

$$\sigma^2 = F\sigma_r^2 + \sigma_Q^2 \quad (1)$$

where σ_r^2 is the formal uncertainty derived from TOA fit-

[1] We implemented the arithmetic using the MLAPACK library: `http://mplapack.sourceforge.net/`.

[2] While formally the VLBI astrometry should be incorporated into the prior in our analysis, it is acceptable to fix the position at the VLBI position because it is of much greater precision, i.e., the Gaussian and delta-function priors are equivalent relative to the precision that can be measured from the data.

ting. The factor F (often referred to as EFAC[3]) modifies σ_R to account for instrumental distortions. The term σ_Q (often referred to as EQUAD) accounts for additional observation-independent uncertainties. Both F and σ_Q are defined independently for each band-backend system.

Red-noise contributions to the TOAs include radio frequency-independent timing noise and dispersion-measure variations. Their contribution is often described by assuming the amplitude of the fluctuations can be described using a power spectrum, which is suitable for wide-sense stationary processes. The simplest model for timing noise we considered was a power-law power spectrum characterised with a spectral index β and an amplitude A:

$$P_{r,\mathrm{PL}}(f) = A\left(\frac{f}{f_{\mathrm{yr}}}\right)^{\beta}, \tag{2}$$

where f_{yr} is a frequency of 1 cycle per year.

In the second model, the power-law is modified to include a spectral flattening

$$P_{r,\mathrm{BL}}(f) = \frac{A(f_c/f_{\mathrm{yr}})^{-\beta}}{\left[1+(f/f_c)^{-\beta/2}\right]^2}, \tag{3}$$

where A again is the amplitude, and f_c is a corner frequency. With this definition, $P_r(f)$ is a power law when $f \gg f_c$, and when $f \ll f_c$, $P_r(f)$ is constant. This model is motivated by observations of non-power-law and apparently quasiperiodic timing noise observed in many pulsars (Hobbs, Lyne & Kramer 2010).

We modelled the noise in the time domain using a harmonically related series of sinusoids constrained to have spectral density parametrized using the forms in Equations (2) or (3). It is necessary to mitigate spectral leakage of low-frequency power when modelling stochastic processes with Fourier series. However if the timing noise is band-limited or the power-law spectral index is relatively flat ($\beta > -6$), it is sufficient to start the series at $f = 1/T_{\mathrm{span}}$, because the inclusion of ν and $\dot{\nu}$ act as a pre-whitening filter on the data set (Blandford, Romani & Narayan 1984). However, for young pulsars, timing noise can be very steep, e.g. PSR B1259−63 was measured to have $\beta \approx -9$ (Shannon, Johnston & Manchester 2014). To mitigate spectral leakage in this case, we included low frequencies $f < 1/T$ to model the lowest frequency timing noise at logarithmically spaced intervals, as described in van Haasteren & Vallisneri (2015).

We also searched for dispersion measure (DM) variations. Following a technique outlined in Lentati et al. (2015), we include two components in the model. The first is a stochastic component, in which the DM fluctuations (distinct from the TOA fluctuations) are modelled as a time series constrained to have a power-law power spectrum,

$$P_{\mathrm{DM}}(f) = D\,(f/f_{\mathrm{yr}})^{\gamma}, \tag{4}$$

where D is the amplitude of the DM variations and γ is the spectral index of the assumed power law. The second component is a quadratic polynomial in DM that accounts for secular trends observed in some pulsars (Keith et al. 2013) but also acts to mitigate spectral leakage of the first component. Together they account for stochastic variations associated with the the turbulent ionised interstellar medium, but the polynomial also models linear trends in DM that are also observed by Keith et al. (2013). In the TOAs, the fluctuations induced by the DM variations scale by the inverse-square of the observing frequency, enabling DM variations and timing noise to be distinguished.

We also searched for pulsar braking, a deterministic process that primarily induces low-frequency signals in TOAs. The pulsar rotation rate is expected to decelerate as the pulsar spins down because of the associated decrease in the magnetic torque. As a pulsar slows down, the rate of deceleration decreases, resulting in a positive second derivative. In general the braking can be approximated by a braking index n which modulates the secular spin evolution of the pulsar:

$$\dot{\nu} = K\nu^n. \tag{5}$$

In standard magnetic-dipole braking, the pulsar magnetic field strength and magnetic-dipole inclination angle are assumed to be constant and $n = 3$. Departures from this value are interpreted as being associated with magnetic field evolution, changes in the spin-magnetic axis orientation, or the effects of a pulsar wind.

For young pulsars this braking index induces a measurable second derivative of the spin frequency $\ddot{\nu}_b$,

$$\ddot{\nu}_b = n\frac{\dot{\nu}^2}{\nu}, \tag{6}$$

and, potentially a third frequency derivative $\dddot{\nu}$,

$$\dddot{\nu}_b = n(2n-1)\frac{\dot{\nu}^3}{\nu^2}. \tag{7}$$

Measured values of n disagree markedly with that predicted from magnetic-dipole braking. For the youngest pulsars ($\tau_c < 10$ kyr) , the braking index generally has a small positive value ($1 < n < 5$). For other young pulsars ($10 < \tau_c < 10^6$ kyr), the braking index is positive, but typically large ($n \gtrsim 10$), suggesting that n is not associated with braking but another process, such as recovery from previous glitches (Johnston & Galloway 1999; Wang et al. 2001; Hobbs, Lyne & Kramer 2010). For the oldest pulsars ($\tau_c > 10^6$ kyr), the braking index is measured to have both positive and negative values, again suggesting it is not associated with pulsar braking but another process. In these pulsars, n is likely being masked by red noise. In previous studies of braking indices, no attempt was made to model the timing noise simultaneously to estimating n. We searched for pulsar braking by including it as a parameter in our timing model and adding in its contribution to $\ddot{\nu}$ and $\dddot{\nu}$ as described in Equations (6) and (7).

4 RESULTS

In panel *a* of Figure 1, we show the residual arrival times, only fitting for the pulsar spin frequency and frequency derivative and therefore excluding terms that account for glitch events or the timing noise. In the absence of these terms, the residual arrival times Δt show peak-to-peak variations of 100 s, which is > 1000 cycles of pulse phase. At the epochs of the glitches (indicated by the stars at the bottom

[3] We have used the TEMPONEST convention for defining EFAC and EQUAD which differs from the TEMPO2 definition.

of the panel), the residuals show a discontinuous change in slope, associated with sudden change in spin frequency. Between glitches, the residuals have positive curvature. This process has been interpreted as a change in $\dot{\nu}$ and has been widely reported (Espinoza et al. 2011; Yu et al. 2013). Interestingly, the process apparently causes the pulsar to return to comparable phase at each epoch, after accounting for the uncertainty in ν and $\dot{\nu}$ through marginalisation.

4.1 Models for glitches

We compare phenomenologically different scenarios for the glitches and timing noise, motivated by previously applied models and the observed residuals presented in panel *a* of Figure 1:

Permanent and transient changes in ν: As a minimal model, we assume that the glitches can be described by a permanent ($\Delta\nu_p$) and transient ($\Delta\nu_t$) change in the spin frequency. The permanent change introduces a change in the pulse phase at time $t > t_g$

$$\Delta\phi(t) = \Delta\nu_p(t - t_g), \tag{8}$$

where t_g is the glitch epoch. The transient component is assumed to exponentially decay on a time scale τ, so its contribution to pulse phase is

$$\Delta\phi(t) = \nu_d\tau\left[1 - \exp\left(\frac{t - t_g}{\tau}\right)\right]. \tag{9}$$

We search over all possible values for the component amplitudes (both positive and negative) and decay times from $1 < \tau < 10^4$ d.

Two decay time scales: We include an additional series of glitch decay times, to search for both long time-scale decays $\Delta\nu_\ell \gg 80$ d and short time-scale glitch decay components $\Delta\nu_s$, which, respectively have decay time-scales τ_ℓ and τ_s.

Common glitch-decay time scales: In addition to assuming that the glitches can be described by permanent and transient changes in spin frequency, we assume that the transient components (long and short) all have the same time scales τ_s and τ_ℓ, respectively.

Permanent changes in $\dot{\nu}$: We assume that there are discrete changes in the spin frequency derivative $\Delta\dot{\nu}_p$ at the epoch of each glitch, as has been previously reported for the pulsar (Yu et al. 2013) and commonly modelled. The change in pulse phase is, for $t > t_g$,

$$\Delta\phi(t) = \frac{1}{2}\Delta\dot{\nu}_p\,(t - t_g)^2 \tag{10}$$

Decay from glitch prior to our observations: We include exponential recovery associated with a glitch that occurred prior to the observation. This would be recovery of a glitch occurring ≈ 400 d prior to our first observation (Flanagan 1991), or to model long time-scale glitch recovery, which, as noted above, is hypothesised to be the source of timing noise in young pulsars.

Additional small glitches: In addition to the reported glitches, we search the dataset for additional glitches. This enables us to assess the apparent dichotomy between glitches and timing noise and search for micro-glitches that have been previously suggested (Cordes, Downs & Krause-Polstorff 1988).

In summary, the general model for a glitch with permanent changes $\Delta\nu_p$, and $\Delta\dot{\nu}_p$, and a single transient component $\Delta\nu_t$ that decays on a timescale τ is

$$\Delta\phi(t) = \Delta\Phi + \Delta\nu_p(t - t_g), + \frac{1}{2}\Delta\dot{\nu}_p\,(t - t_g)^2 + \nu_d\tau\left[1 - \exp\left(\frac{t - t_g}{\tau}\right)\right]. \tag{11}$$

The phase $\Delta\Phi$ accounts for errors in t_g (if it is held fixed) or errors in the pulse numbering at the glitch event.

4.2 Models for timing noise

We consider a smaller number of models for the timing noise:

No timing noise: We first consider a model containing no timing noise. This model therefore assumes that the entirety of the time-correlated signal in the TOAs is associated with the glitches.

Power-law timing noise: We assume that the timing noise can be described by a wide-sense-stationary power-law process (Equation 2) that is observed in many young pulsars.

Band-limited timing noise: We assume that the timing noise is band limited and can be described by Equation (3), which would be the case if it was caused by state changing or a similar process.

Non-stationary timing noise: We assume that in each inter-glitch period, the timing noise can be described by an independent power-law red-noise process. This model enables us to determine if the timing noise and glitch activity are correlated.

Braking index: We additionally search for pulsar braking index n.

4.3 Supported models

For the models listed in the previous sections, we calculated the Bayesian evidence over the distribution of posterior parameters of interest while marginalising over the remaining nuisance parameters.

In Table 1 we show a hierarchy of models that provided significantly improved evidence, culminating in the preferred model. In general, models that provide improvements of $\Delta \log E > 3$ or have fewer parameters but comparable values of evidence are favoured.

1. Power-law stationary timing noise. In panel *b* of Figure 1 we show the residuals from the maximum-likelihood model accounting for the glitches, modelling the timing noise to be a power-law red noise process and marginalising over the unknown ν and $\dot{\nu}$. The dominant signal in the residuals is red noise, which induces peak-to-peak variations in the residuals of 12 s. The residuals are well described by a power-law red noise process with a spectral index of $\beta \approx -6$. We measured a consistent amplitude and spectral shape for the timing noise using the spectral-modelling algorithm presented in Coles et al. (2011), after fixing the glitch parameters at their maximum-likelihood values. We find no evidence that the level or amplitude of timing noise varies between glitches.

2. Permanent frequency changes and two decaying component with time-scales (a short time scale of $\tau_s \approx 25$ d and

Table 1. Model comparison

Model	TN Model	Glitch Model	log(E)	Δlog(E)
1	No red noise	$\Delta\nu_p + \Delta\nu_t$	4559.4	-5394.0
2	Red	$\Delta\nu_p + \Delta\nu_t$	9953.4	0.0
3	Red + DM	$\Delta\nu_p + \Delta\nu_t$	10057.5	104.1
4	Red + DM	$\Delta\nu_p + \Delta\dot{\nu}$	10018.4	65.0
5	Red + DM	(3) + prior glitch	10099.7	146.3
6	Red + DM	(4) + $\Delta\nu_s$ + $\Delta\nu_\ell$	10143.9	190.5
7	Red + DM	(5) + common glitch time scales	10144.7	191.3

Models for timing of the Vela pulsar and their evidence. For each model we show the timing noise model, and the glitch models, and the total evidence, and relative evidence. . Timing noise model: RN: Red noise; DM: dispersion-measure variations. Glitch models: $\Delta\nu_p$: permanent changes in spin frequency; $\Delta\nu_t$: transient change in spin frequency (if only one transient component included), modelled by as an exponential variation with an amplitude and a time scale τ ; $\Delta\nu_\ell$: long duration transient component (if two component are included), with time scale τ_ℓ; $\Delta\nu_s$: short-duration transient component (if two components included) with time scale τ_s; $\Delta\dot{\nu}$: permanent changes in spin-frequency derivative.

a long time scale of $\tau_\ell \approx 1300$ d) that are common amongst the glitches. The amplitude of the glitch components can be found in Table 2. For glitches 1, 2, 3, and 5, our data could not constrain the amplitudes of the short duration components $\Delta\nu_s$, likely because of poor sampling at these epochs. In panel *c* of Figure 1, we show the maximum-likelihood model of the glitch signal, after marginalising over the pulsar spin-down. The signal closely resembles the difference of panels *a* and *b*. Once accounting for the timing noise (and marginalising over the uncertain ν, $\dot{\nu}$, DM, and instrumental jumps) the glitch signal appears to be relatively stationary in the residuals. The maximum-likelihiood parameters for the glitches are displayed in Table 2. We find no evidence for any variation in the glitch-decay time scales. When modelled individually the long glitch decay time-scales show consistent posterior distributions. These long time-scale components have not previously been observed and have been modelled previously as a combination of inter-glitch timing noise and permanent changes in slowdown ($\Delta\dot{\nu}$). Models containing only short time scale glitches and permanent changes in $\dot{\nu}$ (Model 4 in Table 1) are significantly disfavoured over models containing only transient components (Model 5). We also find evidence for the decay from a previous glitch (Model 5 in Table 1). This term accounts for latent exponential decay from a glitch 400 d prior to our observations. In contrast to previous analyses, we find a significant component of each glitch event is associated with a transient component, with between 30% and 80% of the total change in spin frequency associated with a transient component.

3. No small glitches. In addition to modelling the 8 large glitches in our data set, we searched for additional small glitches. We place limits on the amplitudes of other glitches of $\approx 10^{-7}$ averaged over the observing span. Lower amplitude events are covariant with the timing noise.

4. No permanent changes in spin frequency derivative $\Delta\dot{\nu}$. This component instead is attributed to the long time-scale glitch decay. The positive curvature in the residuals between epochs, usually modelled as a change in $\dot{\nu}$ is better modelled by long-term recovery. We verified that our methods could detect differences between permanent changes in $\dot{\nu}$ and long-term decays by simulating data sets that contained either significant $\Delta\dot{\nu}_p$ or long-term decays. The simulated data set had either the maximum-likelihood values of $\Delta\dot{\nu}$ (from a model that did not contain long-term decays) or those from the long time scale decays (as listed in Table 2) obtained from our data set. The simulated data sets also contained the maximum-likelihood realisation of red noise from our data set and simulated TOA uncertainties corrected for EQUAD and EFAC. In both cases the correct model was selected with $\Delta \log E > 130$ (i.e., with probabilities of $> 1 - e^{-130}$).

In panel *d* of Figure 1, we show the *whitened* maximum-likelihood residuals, after accounting for glitches, timing noise, and dispersion-measure variations (discussed below). The residuals show no systematic variations in arrival times, suggesting that the model is complete, and providing further confirmation that the pulse numbering is correct. The plot also shows that the data quality has greatly improved in the most recent data, due almost entirely to instrumental improvements.

4.4 Dispersion-measure variations

As part of a larger study of dispersion-measure (DM) variations young energetic pulsars, Petroff et al. (2013) measured DM variations for the Vela pulsar, utilising some of the data presented here. The DM variations were modelled using a linearly interpolated time series described and as described in Keith et al. (2013). We searched for the amplitude and power-law index associated with DM variations, as described above. The results of our model for DM variations are presented in panel *e* of Figure 1 and are consistent with those presented in Petroff et al. (2013). At 1.4 GHz, the DM variations induce TOA fluctuations that are a factor of 10^4 smaller than the timing noise.

5 DISCUSSION

Compared to previous analyses of the glitches, we find that a much larger fraction of the changes in spin frequency is associated with transient decaying components. The dominant transient components are associated with a long time-scale recovery ($\tau \approx 1300$ d), with a time scale that is common to all glitches. Not surprisingly, the sum of our permanent and

Table 2. Glitch parameters

	MJD	$\Delta\nu_p$ (μHz)	$\Delta\nu_\ell$ (μHz)	$\Delta\nu_s$ (μHz)	$\Delta\nu_g/\nu$ (10^{-6})	Q	$(\Delta\nu_g/\nu)_{\rm lit}$ (10^{-6})	Ref.
P	48457	...	16(4)	...	...	...	2.715(2)	1
1	49559	1.8(2)	7.8(2)	...	0.86(3)	0.81(4)	0.835(2)	2
2	49591	1.4(3)	0.7(2)	...	0.19(2)	0.3(1)	0.199(2)	3
3	50369	12.3(3)	11.5(3)	0.11(3)	2.14(5)	0.48(2)	2.11(2)	4
4	51559	22.7(6)	12.1(6)	0.18(3)	3.12(8)	0.35(2)	3.152(2)	4
5	53193	10.6(3)	12.3(3)	0.1(1)	2.06(4)	0.54(2)	2.100	4
6	53960	19.5(4)	9.2(3)	0.3(1)	2.59(5)	0.83(3)	2.62	4
7	55408	9.1(3)	12.1(3)	...	1.89(4)	0.33(2)	1.94	5
8	56555	20.9(2)	13.2(2)	0.21(2)	3.06(4)	0.39(1)	3.100	6

Maximum-likelihood glitch parameters. For each glitch, we list the MJD, the permanent change in spin frequency $\Delta\nu_p$, the long and short glitch-decay amplitudes (respectively, $\Delta\nu_\ell$ and $\Delta\nu_s$) , reported in absolute value. For glitches 1,2,7, our data could not significantly constrain the amplitudes of the short glitch recoveries. We also show the total change in spin frequency at the glitch epoch $\Delta\nu_g$, which is the sum of the permanent and transient components, relative to the pulsar spin frequency ν. The values in parentheses represent the nominal $1-\sigma$ uncertainties for our measurements and, where available, previous measurements.
$Q = (\Delta\nu_\ell + \Delta\nu_s)/(\Delta\nu_s + \Delta\nu_\ell + \Delta\nu_p)$ is the fraction of the glitch that is recovered. Where available we also show previously reported total glitch-decay amplitudes $\Delta\nu_{g,\rm lit}$. The references are (1) Flanagan (1991); (2) Flanagan & McCulloch (1994); (3) Flanagan (1994) ; (4) Yu et al. (2013) ; (5) Buchner (2010); and (6) Buchner (2013).

transient changes in spin frequency is comparable to previously published measurements of the permanent component as displayed in Table 2. With only eight glitches in our data set, we have an insufficient sample size to determine if there is a correlation between the amplitude of the permanent and transient components to pulsar glitches. The two glitches observed on MJD 49559 and 49591 are significantly smaller in $\Delta\nu_p$ and $\Delta\nu_\ell$ than the other glitches but are also unusual because of their relative contemporaneity.

Many first-principle models of glitches predict that transient components of different glitches should decay with the same time scale (e.g., van Eysden & Melatos 2010), because the microphysics of the neutron star, which regulates glitch decay, does not change.

The length of the decays suggests that the recovery (and the glitches) could be associated with the non-linear regime in the vortex-creep model for glitches. It has been predicted that the transition from linear to non-linear creep regime for the Vela pulsars would occur on time-scales of ≈ 1000 d, close to our glitch time-scale of 1300 d (Alpar et al. 1993). While the regime may be non-linear, the glitch recovery would remain linear because the perturbation is relatively weak (Alpar, Cheng & Pines 1989), with $\Delta\nu_g/|\dot{\nu}| \ll \tau_\ell$.

Most of the inter-glitch TOA variations are associated with the long time-scale decay of the transient component (glitch recovery). This is consistent with observations of timing noise in other young pulsars (Hobbs, Lyne & Kramer 2010), which show $\ddot{\nu} > 0$, much like the Vela pulsar. However, other pulsars show markedly different distributions of glitch waiting times.

We find that the timing noise can be described by a wide-sense stationary process. When we considered independent realisations of timing noise between glitches, we found that they all had consistent amplitudes and spectral indices. Furthermore, the evidence supported a single coherent process. Once accounting for glitch recovery, we exclude long time-scale decay as being the origin of the timing noise. The timing noise has a spectral index of $\beta = -6.0 \pm 0.5$ (1σ), which is comparable to other young pulsars. The stationarity of the timing noise suggests that it is not related to the glitches. Because the timing noise is subdominant and unaffected by the glitches, we have no evidence for any causal or correlated relationship between the two phenomena.

We do not find any evidence for braking of the pulsar spin down and set an upper limit of $n < 8$. Our sensitivity to the braking index is limited by the large TOA variations induced by the glitch events, the presence of timing noise, and the covariance of braking index with these parameters. A previous measurement of braking index relied only on measurements of ν and $\dot{\nu}$ at specific post-glitch epochs to measure the braking index (Lyne et al. 1996). Through 9 glitches, they found measurements of $\dot{\nu}$ at a date ≈ 150 d after each glitch (and after apparent glitch recovery) were consistent with a braking index of 1.4 ± 0.2. We attempted to reproduce the results of Lyne et al. (1996) using our dataset, which span an independent set of glitches. We whitened our dataset by including a series of sinusoids in our maximum-likelihood timing model (displayed in panel b of Figure 1). We then used this model (which is the sum of the glitches, their recovery, and the timing noise) to to calculate $\dot{\nu}$ at epochs 150 d after glitch events. Like Lyne et al. (1996), we find that there is a linear trend in ν at these epochs. However, we do not identify a linear trend in $\dot{\nu}$ at any epoch after the 8 glitches in our observations. We derive an apparent braking index to be 0.0 ± 0.9 at epochs 120 d after glitches. We attribute the differences in the apparent braking index to the the noise realisations in the two independent data sets.

Despite the high levels of rotational instability, the timing behaviour observed in Vela shares similarities to that observed in older pulsars which glitch less and have much lower levels of timing noise. In a sample of 366 pulsars monitored over $\gtrsim 30$ yr, Hobbs, Lyne & Kramer (2010) found a number of pulsars that showed cuspy timing events that occur quasiperiodically. These cuspy events represent increases in ν, followed by episodes of increased $\dot{\nu}$ that are

interrupted (at the same residual phase) by another increase in ν. The best example of this is PSR B1900+06 (Hobbs, Lyne & Kramer 2010). If these cusps are interpreted as small glitches and the intervening episodes as glitch recovery, the rotational irregularities are analogous to what we observe in the Vela pulsar, but with a smaller magnitude.

The robustness of the solution presented here depends on the underlying assumption that one of the model families considered is correct. The greatest uncertainty is the assumption that the timing noise follows a power-law or broken power process. If the timing noise can be better constrained it may be possible to detect and characterise additional components or measure a significant braking index. These uncertainties affect both Bayesian and maximum-likelihood approaches. Searches for the most transient components are limited by the relatively poor (monthly) cadence of our observations.

6 CONCLUSIONS

We have presented a timing solution for the Vela pulsar that spans ≈ 21 yr. The solution is nearly phase connected, with the only uncertainties associated with pulse numbering at glitch events. This solution, use of a full timing model, and Bayesian methodology has enabled the robust parametrization of the glitches and characterise the spin noise.

We have identified dominant transient components to the glitches that decays on a common 25 d and 1300 d time scales for all 8 glitches in our dataset and a subdominant steep red noise component. These methods can be applied to other young pulsars to identify long time-scale glitch decay components, characterise timing noise, and robustly measure braking indices.

ACKNOWLEDGMENTS

The Parkes radio telescope is part of the Australia Telescope National Facility, which is funded by the Commonwealth of Australia for operation as a National Facility managed by CSIRO. GH is the recipient of ARC Future Fellowship FT120100595. LTL acknowledges support from a junior research fellowship at Trinity Hall College, Cambridge University. RMS acknowledges travel support for this work through a John Philip early-career research award from CSIRO.

REFERENCES

Alpar M. A., Chau H. F., Cheng K. S., Pines D., 1993, ApJ, 409, 345
Alpar M. A., Cheng K. S., Pines D., 1989, ApJ, 346, 823
Alpar M. A., Pines D., Anderson P. W., Shaham J., 1984, ApJ, 276, 325
Anderson P. W., Itoh N., 1975, Nature, 256, 25
Baym G., Pethick C., Pines D., 1969, Nature, 224, 673
Blandford R., Romani R. W., Narayan R., 1984, Journal of Astrophysics and Astronomy, 5, 369
Buchner S., 2013, The Astronomer's Telegram, 5406, 1
Buchner S. J., 2010, The Astronomer's Telegram, 2768
Coles W., Hobbs G., Champion D. J., Manchester R. N., Verbiest J. P. W., 2011, MNRAS, 418, 561
Cordes J. M., Downs G. S., 1985, ApJS, 59, 343
Cordes J. M., Downs G. S., Krause-Polstorff J., 1988, ApJ, 330, 847
Cordes J. M., Helfand D. J., 1980, ApJ, 239, 640
Dodson R., Legge D., Reynolds J. E., McCulloch P. M., 2003, ApJ, 596, 1137
Dodson R. G., McCulloch P. M., Lewis D. R., 2002, ApJL, 564, L85
Edwards R. T., Hobbs G. B., Manchester R. N., 2006, MNRAS, 372, 1549
Espinoza C. M., Lyne A. G., Stappers B. W., Kramer M., 2011, MNRAS, 414, 1679
Feroz F., Hobson M. P., Bridges M., 2009, MNRAS, 398, 1601
Flanagan C., 1991, IAU circ., 5311, 3
Flanagan C., 1994, IAU circ., 6064, 2
Flanagan C., McCulloch P. M., 1994, IAU circ., 6038, 2
Groth E. J., 1975, ApJS, 29, 453
Handley W. J., Hobson M. P., Lasenby A. N., 2015, arXiv:1502.01856
Haskell B., Melatos A., 2015, International Journal of Modern Physics D, 24, 30008
Hobbs G., Lorimer D. R., Lyne A. G., Kramer M., 2005, MNRAS, 360, 974
Hobbs G., Lyne A. G., Kramer M., 2010, MNRAS, 402, 1027
Hobbs G. B., Edwards R. T., Manchester R. N., 2006, MNRAS, 369, 655
Hotan A. W., van Straten W., Manchester R. N., 2004, PASA, 21, 302
Jenet F. A., Anderson S. B., Kaspi V. M., Prince T. A., Unwin S. C., 1998, ApJ, 498, 365
Johnston S., Galloway D., 1999, MNRAS, 306, L50
Keith M. J. et al., 2013, MNRAS, 429, 2161
Kerr M., Johnston S., Hobbs G., Shannon R. M., 2015, arXiv:1507.06982
Large M. I., Vaughan A. E., Mills B. Y., 1968, Nature, 220, 340
Lentati L., Alexander P., Hobson M. P., Feroz F., van Haasteren R., Lee K. J., Shannon R. M., 2014, MNRAS, 437, 3004
Lentati L., Shannon R. M., 2015, MNRAS, 454, 1058
Lentati L. et al., 2015, arXiv:1504.03692
Livingstone M. A., Ng C.-Y., Kaspi V. M., Gavriil F. P., Gotthelf E. V., 2011, ApJ, 730, 66
Lyne A., Hobbs G., Kramer M., Stairs I., Stappers B., 2010, Science, 329, 408
Lyne A. G., Jordan C. A., Graham-Smith F., Espinoza C. M., Stappers B. W., Weltevrede P., 2015, MNRAS, 446, 857
Lyne A. G., Pritchard R. S., Graham-Smith F., Camilo F., 1996, Nature, 381, 497
Manchester R. N., Hobbs G., 2011, ApJL, 736, L31
Melatos A., Link B., 2014, MNRAS, 437, 21
Melatos A., Warszawski L., 2009, ApJ, 700, 1524
Petroff E., Keith M. J., Johnston S., van Straten W., Shannon R. M., 2013, MNRAS, 435, 1610
Radhakrishnan V., Manchester R. N., 1969, Nature, 222, 228
Reichley P. E., Downs G. S., 1969, Nature, 222, 229

Ruderman M., 1969, Nature, 223, 597
Shannon R. M., Cordes J. M., 2010, ApJ, 725, 1607
Shannon R. M. et al., 2013, ApJ, 766, 5
Shannon R. M., Johnston S., Manchester R. N., 2014, MNRAS, 437, 3255
Shannon R. M. et al., 2015, Science, 349, 1522
Taylor J. H., 1992, Royal Society of London Philosophical Transactions Series A, 341, 117
van Eysden C. A., Melatos A., 2010, MNRAS, 409, 1253
van Haasteren R., Levin Y., 2013, MNRAS, 428, 1147
van Haasteren R., Levin Y., McDonald P., Lu T., 2009, MNRAS, 395, 1005
van Haasteren R., Vallisneri M., 2015, MNRAS, 446, 1170
Wang N., Manchester R. N., Pace R. T., Bailes M., Kaspi V. M., Stappers B. W., Lyne A. G., 2000, MNRAS, 317, 843
Wang N., Manchester R. N., Zhang J., Wu X. J., Yusup A., Lyne A. G., Cheng K. S., Chen M. Z., 2001, MNRAS, 328, 855
Weltevrede P. et al., 2010, PASA, 27, 64
Yu M. et al., 2013, MNRAS, 429, 688